arXiv:1701.07208v1 [cs.DS] 25 Jan 2017

A Quasi-Polynomial Approximation for the Restricted Assignment Problem*

Klaus Jansen Lars Rohwedder

Department of Computer Science, University of Kiel, 24118 Kiel, Germany
{kj, lro}@informatik.uni-kiel.de

Abstract

Scheduling jobs on unrelated machines and minimizing the makespan is a classical problem in combinatorial optimization. In this problem a job j has a processing time p_{ij} for every machine i. The best polynomial algorithm known goes back to Lenstra et al. and has an approximation ratio of 2. In this paper we study the Restricted Assignment problem, which is the special case where $p_{ij} \in \{p_j, \infty\}$. We present an algorithm for this problem with an approximation ratio of $11/6 + \epsilon$ and quasi-polynomial running time $n^{\mathcal{O}(1/\epsilon \log(n))}$ for every $\epsilon > 0$. This closes the gap to the best estimation algorithm known for the problem with regard to quasi-polynomial running time.

Keywords: approximation, scheduling, unrelated machines, local search

1 Introduction

In the problem we consider, which is known as Scheduling on Unrelated Machines, a schedule $\sigma : \mathcal{J} \to \mathcal{M}$ of the jobs $\mathcal{J}$ to the machines $\mathcal{M}$ has to be computed. On machine i the job j has a processing time of p_{ij}. We want to minimize the makespan, i.e., $\max_{i \in \mathcal{M}} \sum_{j \in \sigma^{-1}(i)} p_{ij}$. The classical 2-approximation by Lenstra et al. [8] is still the algorithm of choice for this problem.

Recently a special case, namely the Restricted Assignment problem, has drawn much attention in the scientific community. Here each job j has a processing time p_j, which is independent from the machines, and a set of machines $\Gamma(j)$. A job j can only be assigned to $\Gamma(j)$. This is equivalent to the former problem when $p_{ij} \in \{p_j, \infty\}$. For both the general and the restricted variant there cannot be a polynomial algorithm with an approximation ratio better than $3/2$, unless P = NP [8]. If the exponential time hypothesis (ETH) holds, such an algorithm does not even exist with sub-exponential (in particular, quasi-polynomial) running time [5].

In a recent breakthrough, Svensson has proved that the configuration-LP, a natural linear programming relaxation, has an integrality gap of at most $33/17$ [10]. We have later improved this bound to $11/6$ [7]. By approximating the configuration-LP this yields an $(11/6+\epsilon)$-estimation algorithm for every $\epsilon > 0$. However, no polynomial algorithm is known that can produce a solution of this value.

For instances with only two processing times additional progress has been made. Chakrabarty et al. gave a polynomial $(2 - \delta)$-approximation for a very small δ [4]. Later Annamalai surpassed this with a $(17/9 + \epsilon)$-approximation for every $\epsilon > 0$ [1]. For this special case it was also shown that the integrality gap is at most $5/3$ [6].

In [10], [7], and [6] the critical idea is to design a local search algorithm, which is then shown to produce good solutions. However, the algorithm has a potentially high running time; so it was only used to prove the existence of such a solution. A similar algorithm was used in the Restricted Max-Min Fair Allocation problem. Here a quasi-polynomial variant by Polácek et al. [9] and a polynomial variant by Annamalai et al. [2] were later discovered.

*Research supported by German Research Foundation (DFG) project JA 612/15-1

$$\sum_{C \in \mathcal{C}_i(T)} x_{i,C} \leq 1 \qquad \forall i \in \mathcal{M}$$

$$\sum_{i \in \mathcal{M}} \sum_{C \in \mathcal{C}_i(T): j \in C} x_{i,C} \geq 1 \qquad \forall j \in \mathcal{J}$$

$$x_{i,C} \geq 0$$

$$\min \sum_{i \in \mathcal{M}} y_i - \sum_{j \in \mathcal{J}} z_j$$

$$s.t.$$

$$y_i \geq \sum_{j \in C} z_j \quad \forall i \in \mathcal{M}, C \in \mathcal{C}_i(T)$$

$$y_i, z_j \geq 0$$

Figure 1: Primal (left) and dual (right) of the configuration-LP w.r.t. makespan T

In this paper, we present a variant of the local search algorithm, that admits a quasi-polynomial running time. The algorithm is purely combinatorial and uses the configuration-LP only in the analysis.

Theorem 1.1. *For every $\epsilon > 0$ there is an $(11/6 + \epsilon)$-approximation algorithm for the* Restricted Assignment *problem with running time $\exp(\mathcal{O}(1/\epsilon \cdot \log^2(n)))$, where $n = |\mathcal{J}| + |\mathcal{M}|$.*

The main idea is the concept of layers. The central data structure in the local search algorithm is a tree of so-called blockers and we partition this tree into layers, that are closely related to the distance of a blocker from the root. Roughly speaking, we prevent the tree from growing arbitrarily high. A similar approach was taken in [9].

1.1 The configuration-LP

A well known relaxation for the problem of Scheduling on Unrelated Machines is the configuration-LP (see Fig. 1). The set of configurations with respect to a makespan T are defined as $\mathcal{C}_i(T) = \{C \subseteq \mathcal{J} : \sum_{j \in C} p_{ij} \leq T\}$. We refer to the minimal T for which this LP is feasible as the optimum or OPT^*. In the Restricted Assignment problem a job j can only be used in configurations of machines in $\Gamma(j)$ given T is finite. We can find a solution for the LP with a value of at most $(1+\epsilon)\mathrm{OPT}^*$ in polynomial time for every $\epsilon > 0$ [3].

1.2 Preliminaries

In this section we simplify the problem we need to solve. The approximation ratio we will aim for is $1 + R$, where $R = 5/6 + 2\epsilon$. We assume that $\epsilon < 1/12$ for our algorithm, since otherwise the 2-approximation in [8] can be used.

We will use a binary search to obtain a guess T for the value of OPT^*. In each iteration, our algorithm either returns a schedule with makespan at most $(1 + R)T$ or proves that T is smaller than OPT^*. After polynomially many iterations, we will have a solution with makespan at most $(1 + R)\mathrm{OPT}^*$. To shorten notation, we scale each size by $1/T$ within an iteration, that is to say our algorithm has to find a schedule of makespan $1 + R$ or show that $\mathrm{OPT}^* > 1$. Unless otherwise stated we will assume that $T = 1$ when speaking about configurations or feasibility of the configuration-LP.

Definition 1.2 (Small, big, medium, huge jobs)**.** A job j is small if $p_j \leq 1/2$ and big otherwise; A big job is medium if $p_j \leq 5/6$ and huge if $p_j > 5/6$.

The sets of small (big, medium, huge) jobs are denoted by $\mathcal{J}_S$ (respectively, $\mathcal{J}_B$, $\mathcal{J}_M$, $\mathcal{J}_H$). Note that at most one big job can be in a configuration (w.r.t. makespan 1).

Definition 1.3 (Valid partial schedule)**.** We call $\sigma : \mathcal{J} \to \mathcal{M} \cup \{\perp\}$ a valid partial schedule if (1) for each job j we have $\sigma(j) \in \Gamma(j) \cup \{\perp\}$, (2) for each machine $i \in \mathcal{M}$ we have $p(\sigma^{-1}(i)) \leq 1 + R$, and (3) each machine is assigned at most one huge job.

$\sigma(j) = \perp$ means that job j has not been assigned. In each iteration of the binary search, we will first find a valid partial schedule for all medium and small jobs and then extend the schedule one huge job at

a time. We can find a schedule for all small and medium jobs with makespan at most 11/6 by applying the algorithm by Lenstra, Shmoys, and Tardos [8]. This algorithm outputs a solution with makespan at most $\mathrm{OPT}^* + p_{max}$, where p_{max} is the biggest processing time (in our case at most 5/6). The problem that remains to be solved is given in below.

Input: An instance of RESTRICTED ASSIGNMENT, a valid partial schedule σ, a huge job j_{new} with $\sigma(j_{\text{new}}) = \perp$.

Output: Either: (1) A valid partial schedule σ' with $\sigma'(j_{\text{new}}) \neq \perp$ and $\sigma(j) \neq \perp \Rightarrow \sigma'(j) \neq \perp$ for all $j \in \mathcal{J}$, or (2) 'error' (indicating that $\mathrm{OPT}^* > 1$).

Without loss of generality let us assume that the jobs are identified by natural numbers, that is $\mathcal{J} = \{1, 2, \ldots, |\mathcal{J}|\}$, and $p_1 \leq p_2 \leq \ldots \leq p_{|\mathcal{J}|}$. This gives us a total order on the jobs that will simplify the algorithm.

2 Algorithm

Throughout the paper, we make use of modified processing times $\overline{P}_j$ and $\overline{p}_j$, which we obtain by rounding the sizes of huge jobs up or down, that is

$$\overline{P}_j = \begin{cases} 1 & \text{if } p_j > 5/6, \\ p_j & \text{if } p_j \leq 5/6; \end{cases} \quad \text{and} \quad \overline{p}_j = \begin{cases} 5/6 & \text{if } p_j > 5/6, \\ p_j & \text{if } p_j \leq 5/6. \end{cases}$$

Definition 2.1 (Moves, valid moves)**.** A pair (j, i) of a job j and a machine i is a *move*, if $i \in \Gamma(j) \backslash \{\sigma(j)\}$. A move (j, i) is *valid*, if (1) $\overline{P}(\sigma^{-1}(i)) + p_j \leq 1 + R$ and (2) j is not huge or no huge job is already on i.

We note that by performing a valid move (j, i) the properties of a valid partial schedule are not compromised.

Definition 2.2 (Blockers)**.** A blocker is a tuple (j, i, Θ), where (j, i) is a move and Θ is the type of the blocker. There are 6 types with the following abbreviations: (SA) *small-to-any blockers*, (HA) *huge-to-any blockers*, (MA) *medium-to-any blockers*, (BH) *huge-/medium-to-huge blockers*, (HM) *huge-to-medium blockers*, and (HL) *huge-to-least blockers*.

The algorithm maintains a set of blockers called the blocker tree $\mathcal{T}$. We will discuss the tree analogy later. The blockers wrap moves that the algorithm would like to execute. By abuse of notation, we write that a move (j, i) is in $\mathcal{T}$, if there is a blocker (j, i, Θ) in $\mathcal{T}$ for some Θ. The type Θ determines how the algorithm treats the machine i as we will elaborate below.

The first part of a type's name refers to the size of the blocker's job, e.g., small-to-any blockers are only used with small jobs, huge-to-any blockers only with huge jobs, etc. The latter part of the type's name describes the undesirable jobs on the machine: The algorithm will try to remove jobs from this machine if they are undesirable; at the same time it does not attempt to add such jobs to the machine. On machines of small-/medium-/huge-to-any blockers all jobs are undesirable; on machines of huge-/medium-to-huge blockers huge jobs are undesirable; on machines of huge-to-medium blockers medium jobs are undesirable and finally on machines of huge-to-least blockers only those medium jobs of index smaller or equal to the smallest medium job on i are undesirable.

The same machine can appear more than once in the blocker tree. In that case, the undesirable jobs are the union of the undesirable jobs from all types. Also, the same job can appear multiple times in different blockers.

The blockers corresponding to specific types are written as $\mathcal{T}_{SA}$, $\mathcal{T}_{HA}$, etc. From the blocker tree, we derive the machine set $\mathcal{M}(\mathcal{T})$ which consists of all machines corresponding to moves in $\mathcal{T}$. This notation is also used with subsets of $\mathcal{T}$, e.g., $\mathcal{M}(\mathcal{T}_{HA})$.

Definition 2.3 (Blocked small jobs, active jobs)**.** A small job j is *blocked*, if it is undesirable on all other machines it it allowed on, that is $\Gamma(j) \backslash \{\sigma(j)\} \subseteq \mathcal{M}(\mathcal{T}_{SA} \cup \mathcal{T}_{MA} \cup \mathcal{T}_{HA})$. We denote the set of blocked small jobs by $S(\mathcal{T})$. The set of active jobs $\mathcal{A}$ includes j_{new}, $S(\mathcal{T})$ as well as all those jobs, that are undesirable on the machine, they are currently assigned to.

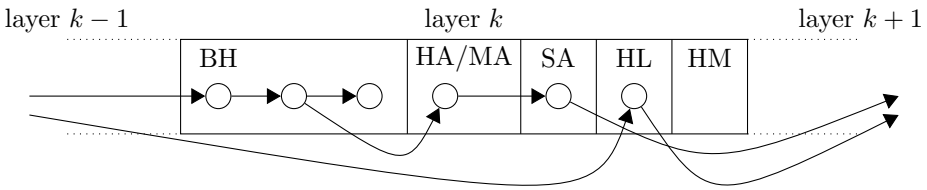

Figure 2: Example layer

We define for all machines i the job sets $S_i(\mathcal{T}) = S(\mathcal{T}) \cap \sigma^{-1}(i)$, $\mathcal{A}_i(\mathcal{T}) = \mathcal{A}(\mathcal{T}) \cap \sigma^{-1}(i)$, $M_i = \sigma^{-1}(i) \cap \mathcal{J}_M$ and $H_i = \sigma^{-1}(i) \cap \mathcal{J}_H$. Moreover, set $M_i^{\min} = \{\min M_i\}$ if $M_i \neq \emptyset$ and $M_i^{\min} = \emptyset$ otherwise.

2.1 Tree and layers

The blockers in $\mathcal{T}$ and an additional root can be imagined as a tree. The parent of each blocker $\mathcal{B} = (j, i, \Theta)$ is only determined by j. If $j = j_{\text{new}}$ it is the root node; otherwise it is a blocker $\mathcal{B}' \in \mathcal{T}$ for machine $\sigma(j)$ with a type for which j is regarded undesirable. If this applies to several blockers, we use the one that was added to the blocker tree first. We say that $\mathcal{B}'$ *activates* j.

Let us now introduce the notion of a layer. Each blocker is assigned to exactly one layer. The layer roughly correlates with the distance of the blocker to the root node. In this sense, the children of a blocker are usually in the next layer. There are some exceptions, however, in which a child is in the same layer as its parent. We now define the layer of the children of a blocker $\mathcal{B}$ in layer k.

1. If $\mathcal{B}$ is a huge-/medium-to-huge blocker, all its children are in layer k as well;
2. if $\mathcal{B}$ is a huge-to-any blocker, children regarding small jobs are in layer k as well;
3. in every other case, the children are in layer $k+1$.

We note that by this definition for an active job j all blockers $(j, i, \Theta) \in \mathcal{T}$ must be in the same layer; in other words, it is unambiguous in which layer blockers for it would be placed in. We say j is *k-headed*, if blockers for j would be placed in layer k. The blockers in layer k are denoted by $\mathcal{T}^{(k)}$. The set of blockers in layer k and below is referred to by $\mathcal{T}^{(\leq k)}$. We use this notation in combination with qualifiers for the type of blocker, e.g., $\mathcal{T}_{HA}^{(k)}$.

We establish an order between the types of blockers within a layer and refer to this order as the sublayer number. The huge-/medium-to-huge blockers form the first sublayer of each layer, huge-to-any and medium-to-any blockers the second, small-to-any blockers the third, huge-to-least the fourth and huge-to-medium blockers the fifth sublayer (see also Table 1 and Figure 2). By saying a sublayer is after (before) another sublayer we mean that either its layer is higher (lower) or both layers are the same and its sublayer number is higher (lower).

In the final algorithm whenever we remove one blocker, we also remove all blockers in its sublayer and all later sublayers (in particular, all descendants). Also, when we add a blocker to a sublayer, we remove all later sublayers. Among other properties, this guarantees that the connectivity of the tree is never compromised. It also means that, if j is undesirable regarding several blockers for $\sigma(j)$, then the parent is in the lowest sublayer among these blockers, since a blocker in a lower sublayer cannot have been added after one in a higher sublayer.

The running time will be exponential in the number of layers; hence this should be fairly small. We introduce an upper bound $K = 2/\epsilon\lceil\ln(|\mathcal{M}|) + 1\rceil = \mathcal{O}(1/\epsilon \cdot \log(|\mathcal{M}|))$ and will not add any blockers to a layer higher than K.

2.2 Detailed description of the algorithm

The algorithm (see Algorithm 1) contains a loop that terminates once j_{new} is assigned. In each iteration the algorithm performs a valid move in the blocker tree if possible and otherwise adds a new blocker.

Algorithm 1: Quasi-polynomial local search

```
1  initialize empty blocker tree 𝒯;
2  loop
3    if a move in 𝒯 is valid then
4      choose a blocker (j, i, Θ) in the lowest sublayer,
5          where (j, i) is valid;
6      let ℬ be the blocker that activated j;
7      // Update the schedule
8      σ(j) ← i;
9      remove all sublayers after ℬ from 𝒯;
10     if j = j_new then
11       return σ;
12     end
13     if not conditions*(ℬ) then
14       remove the sublayer of ℬ from 𝒯;
15     end
16   else
17     let ℓ be the minimum layer to which we can
18         add a potential move;
19     if ℓ > K or no such ℓ exists then
20       return 'error';
21     end
22     add potential move (j, i) of highest priority to layer ℓ;
23     remove all sublayers after (j, i) from 𝒯;
24   end
25 end
```

Adding blockers. We only add a move to $\mathcal{T}$, if it meets certain requirements. A move that does is called a potential move. For each type of blocker we also define a type of potential move: Potential small-to-any moves, potential huge-to-any moves, etc. When a potential move is added to the blocker tree, its type will then be used for the blocker. Let k be a layer and let $j \in \mathcal{A}(\mathcal{T})$ be k-headed. For a move (j, i) to be a potential move of a certain type, it has to meet the following requirements.

1. (j, i) is not already in $\mathcal{T}$;
2. the size of j corresponds to the type, for instance, if j is big, (j, i) cannot be a small-to-any move;
3. j is not undesirable on i w.r.t. $\mathcal{T}^{(\leq k)}$, i.e., (a) $i \notin \mathcal{M}(\mathcal{T}_{SA}^{(\leq k)} \cup \mathcal{T}_{MA}^{(\leq k)} \cup \mathcal{T}_{HA}^{(\leq k)})$ and (b) if j is huge, then $i \notin \mathcal{M}(\mathcal{T}_{BH}^{(\leq k)})$; (c) if j is medium, then $i \notin \mathcal{M}(\mathcal{T}_{HM}^{(\leq k)})$ and either $i \notin \mathcal{M}(\mathcal{T}_{HL}^{(\leq k)})$ or $\min M_i < j$.
4. The load of the target machine has to meet certain conditions (see Table 1).

Comparing the conditions in the table we notice that for moves of small and medium jobs there is always exactly one type that applies. For huge jobs there is exactly one type if $p(S_i(\mathcal{T}^{(\leq k)})) + p_j \leq 1 + R$ and no type applies, if $p(S_i(\mathcal{T}^{(\leq k)})) + p_j > 1 + R$. The table also lists a priority for each type of move. It is worth mentioning that the priority does not directly correlate with the sublayer. The algorithm will choose the move that can be added to the lowest layer and among those has the highest priority. After adding a blocker, all higher sublayers are deleted.

Performing valid moves. The algorithm performs a valid move in $\mathcal{T}$ if there is one. It chooses a blocker (j, i, Θ) in $\mathcal{T}$, where the blocker's sublayer is minimal and (j, i) is valid. Besides assigning j to i, $\mathcal{T}$ has to be updated as well.

Let $\mathcal{B}$ be the blocker that activated j. When certain conditions for $\mathcal{B}$ are no longer met, we will delete $\mathcal{B}$ and its sublayer. The conditions that need to be checked depend on the type of $\mathcal{B}$ and are marked in Table 1 with a star $(*)$. In any case, the algorithm will discards all blockers in higher sublayers than $\mathcal{B}$ is.

Table 1: Types of blockers/potential moves

type		Conditions	S	P	Undesirable
Huge-/ Medium-to-huge (BH)		$p(\sigma^{-1}(i)\backslash H_i) + p_j \leq 1 + R$	1	5	Huge jobs
Small-to-any (SA)		None	3	4	All jobs
Medium-to-any (MA)	$*$	$p(\sigma^{-1}(i)\backslash H_i) + p_j > 1 + R$	2	3	All jobs
Huge-to-any (HA)	$*$	$p(\sigma^{-1}(i)\backslash H_i) + p_j > 1 + R$ $p(S_i(\mathcal{T}^{(\leq k)}) \cup M_i) + p_j \leq 1 + R$	2	3	All jobs
Huge-to-least (HL)	$*$	$p(S_i(\mathcal{T}^{(\leq k)}) \cup M_i^{\min}) + p_j > 1 + R$ $p(S_i(\mathcal{T}^{(\leq k)})) + p_j \leq 1 + R$	4	2	Medium jobs j_M with $j_M \leq \min M_i$
Huge-to-medium (HM)	$*$ $*$	$p(S_i(\mathcal{T}^{(\leq k)}) \cup M_i) + p_j > 1 + R$ $p(S_i(\mathcal{T}^{(\leq k)}) \cup M_i^{\min}) + p_j \leq 1 + R$	5	1	Medium jobs

The conditions are meant in respect to a move (j, i) where j is k-headed. Column S stands for the sublayer and P for the priority of a blocker type. Conditions marked with a star $(*)$ are additionally checked whenever a job activated by this blocker is moved.

3 Analysis

The analysis of the algorithm has two critical parts. First, we show that it does not get stuck, i.e., there is always a blocker that can be added to the blocker tree or a move that can be executed. Then we show that the number of iterations is bounded by $\exp(\mathcal{O}(1/\epsilon \log^2(n)))$.

Theorem 3.1. *If the algorithm returns 'error', then* $\mathrm{OPT}^* > 1$.

The proof consists in the construction of a solution (z^*, y^*) for the dual of the configuration-LP. The value z_j^* is composed of $\overline{p}_j$ and a scaling coefficient (a power of $\delta := 1 - \epsilon$). The idea of the scaling coefficient is that values for jobs activated in higher layers are supposed to get smaller and smaller. We set $z_j^* = 0$ if $j \notin \mathcal{A}(\mathcal{T})$ and $z_j^* = \delta^k \cdot \overline{p}_j$, if $j \in \mathcal{A}(\mathcal{T})$ and k is the smallest layer such that j is k-headed or $j \in S(\mathcal{T}^{(\leq k)})$. For all $i \in \mathcal{M}$ let

$$w_i = \begin{cases} z^*(\mathcal{A}_i(\mathcal{T})) + \delta^k \frac{1}{6} & \text{if } i \in \mathcal{M}(\mathcal{T}_{HA}^{(k)}), \\ z^*(\mathcal{A}_i(\mathcal{T})) - \delta^k \frac{1}{6} & \text{if } i \in \mathcal{M}(\mathcal{T}_{SA}^{(k)}), \\ z^*(\mathcal{A}_i(\mathcal{T})) & \text{otherwise.} \end{cases}$$

Finally set $y_i^* = \delta^K + w_i$. Note that w is well-defined, since a machine i can be in at most one of the sets $\mathcal{M}(\mathcal{T}_{HA}^{(1)}), \mathcal{M}(\mathcal{T}_{SA}^{(1)}), \mathcal{M}(\mathcal{T}_{HA}^{(2)}), \mathcal{M}(\mathcal{T}_{SA}^{(2)}), \ldots$ On a small-/huge-to-any blocker all jobs are undesirable, that is to say as long as one of such blockers remains in the blocker tree, the algorithm will not add another blocker with the same machine. Also note that $z^*(\mathcal{A}_i(\mathcal{T}))$ and $z^*(\sigma^{-1}(i))$ are interchangeable.

Lemma 3.2. *If there is no valid move in $\mathcal{T}$ and no potential move of a k-headed job for a $k \leq K$, the value of the solution is negative, i.e., $\sum_{j \in \mathcal{J}} z_j^* > \sum_{i \in \mathcal{M}} y_i^*$.*

Proof. Using the Taylor series and $\epsilon < 1/12$ it is easy to check $\ln(1 - \epsilon) \geq -\epsilon/2$. This gives

$$K \geq \frac{2}{\epsilon}(\ln(|\mathcal{M}|) + 1) \geq \frac{\ln(2|\mathcal{M}|)}{\epsilon/2} \geq -\frac{\ln(2|\mathcal{M}|)}{\ln(1 - \epsilon)} = \log_\delta \left(\frac{1}{2|\mathcal{M}|} \right).$$

Claim 1. *(see appendix B). For all $k \leq K$ we have $|\mathcal{M}(\mathcal{T}_{HA}^{(k)})| \leq |\mathcal{M}(\mathcal{T}_{SA}^{(k)})|$.*

Using this claim we find that

$$\begin{aligned}\sum_{j\in\mathcal{J}} z^*_j &\geq z^*_{j_{new}} + \sum_{i\in\mathcal{M}} z^*(\sigma^{-1}(i)) \\ &\geq \delta^1\frac{5}{6} + \sum_{i\in\mathcal{M}} y^*_i - \delta^K|\mathcal{M}| + \sum_{k=1}^{K}[\delta^k\frac{1}{6}|\mathcal{M}(\mathcal{T}^{(k)}_{SA})| - \delta^k\frac{1}{6}|\mathcal{M}(\mathcal{T}^{(k)}_{HA})|] \\ &\geq \delta^1\frac{5}{6} + \sum_{i\in\mathcal{M}} y^*_i - \frac{1}{2} + 0 > \sum_{i\in\mathcal{M}} y^*_i.\end{aligned}$$

□

Lemma 3.3. *If there is no valid move in $\mathcal{T}$ and no potential move of a k-headed job for a $k \leq K$, the solution is feasible, i.e., $z^*(C) \leq y^*_i$ for all $i \in \mathcal{M}$, $C \in \mathcal{C}_i$.*

Proof. We will make the following assumptions, that are proved in the appendix with an exhaustive case analysis.

Claim 2. *(see appendix C). Let $k \leq K$, $i \notin \mathcal{M}(\mathcal{T}^{(\leq k)}_{SA} \cup \mathcal{T}^{(\leq k)}_{MA} \cup \mathcal{T}^{(\leq k)}_{HA})$, $C \in \mathcal{C}_i$, $j \in C$ k-headed and big with $\sigma(j) \neq i$. Then $z^*_j \leq z^*(\mathcal{A}_i(\mathcal{T}^{(\leq k)})\backslash C)$.*

Claim 3. *(see appendix C). Let $k \leq K$ and $i \in \mathcal{M}(\mathcal{T}^{(k)}_{SA} \cup \mathcal{T}^{(k)}_{MA} \cup \mathcal{T}^{(k)}_{HA})$. Then*

$$w_i \geq z^*(\mathcal{A}_i(\mathcal{T})) + \delta^k \cdot (1 - \delta\overline{p}(\mathcal{A}_i(\mathcal{T}))).$$

Let $C_0 \in \mathcal{C}_i$ and $C \subseteq C_0$ denote the set of jobs j with $z^*_j \geq \delta^K\overline{p}_j$. In particular, C does not contain jobs that have potential moves. It is sufficient to show that $z^*(C) \leq w_i$, as this would imply

$$z^*(C_0) = z^*(C) + z^*(C_0\backslash C) \leq w_i + \delta^K\overline{p}(C_0) \leq y^*_i.$$

Loosely speaking, the purpose of δ^K in the definition of y^* is to compensate for ignoring all $(K+1)$-headed jobs.

First, consider the case where $i \notin \mathcal{M}(\mathcal{T}_{SA} \cup \mathcal{T}_{MA} \cup \mathcal{T}_{HA})$. There cannot be a small and activated job $j_S \in C$ with $\sigma(j_S) \neq i$, because then (j_S, i) would be a potential move; hence $C \cap \mathcal{J}_S \cap \mathcal{A}(\mathcal{T}) \subseteq C \cap \mathcal{A}_i(\mathcal{T})$. If there is a big job $j_B \in C$ with $\sigma(j_B) \neq i$, then

$$z^*(C) = z^*_{j_B} + z^*(C \cap \mathcal{J}_S) \leq z^*(\mathcal{A}_i(\mathcal{T})\backslash C) + z^*(C \cap \mathcal{A}_i(\mathcal{T})) = z^*(\mathcal{A}_i(\mathcal{T})) = w_i.$$

If there is no such job, then $C \cap \mathcal{A}(\mathcal{T}) \subseteq \mathcal{A}_i(\mathcal{T})$ and in particular $z^*(C) \leq w_i$.

In the remainder of this proof we assume that $i \in \mathcal{M}(\mathcal{T}^{(\ell+1)}_{SA} \cup \mathcal{T}^{(\ell+1)}_{MA} \cup \mathcal{T}^{(\ell+1)}_{HA})$. Note that for any $k \neq \ell+1$ we have $i \notin \mathcal{M}(\mathcal{T}^{(k)}_{SA} \cup \mathcal{T}^{(k)}_{MA} \cup \mathcal{T}^{(k)}_{HA})$. Also, since all jobs on i are active we have that $z^*_j \geq \delta^{\ell+2}\overline{p}_j$ for all $j \in \sigma^{-1}(i)$. Because there is no potential move (j_S, i) for a small job j_S with $z^*_{j_S} \geq \delta^\ell\overline{p}_{j_S}$, we have for all small jobs $j_S \in C\backslash\mathcal{A}_i(\mathcal{T})$: $z^*_{j_S} \leq \delta^{\ell+1}\overline{p}_{j_S}$.

Case 1. **For every big job $j \in C$ with $\sigma(j) \neq i$ we have $z^*_j \leq \delta^{\ell+1}\overline{p}_j$.** This implies

$$\begin{aligned}z^*(C\backslash\mathcal{A}_i(\mathcal{T})) \leq \delta^{\ell+1}\overline{p}(C\backslash\mathcal{A}_i(\mathcal{T})) &= \delta^{\ell+1}(\overline{p}(C) - \overline{p}(\mathcal{A}_i(\mathcal{T}) \cap C)) \\ &\leq \delta^{\ell+1}(1 - \delta\overline{p}(\mathcal{A}_i(\mathcal{T}) \cap C)).\end{aligned}$$

Therefore

$$\begin{aligned}z^*(C) &= z^*(\mathcal{A}_i(\mathcal{T}) \cap C) + z^*(C\backslash\mathcal{A}_i(\mathcal{T})) \\ &\leq z^*(\mathcal{A}_i(\mathcal{T}) \cap C) + \delta^{\ell+1}(1 - \delta\overline{p}(\mathcal{A}_i(\mathcal{T}) \cap C)) \\ &\leq z^*(\mathcal{A}_i(\mathcal{T})) + \delta^{\ell+1}(1 - \delta\overline{p}(\mathcal{A}_i(\mathcal{T}))) \leq w_i.\end{aligned}$$

Case 2. **There is a big job $j \in C$ with $\sigma(j) \neq i$ and $z_j^* \geq \delta^\ell \overline{p}_j$.** Let $k \leq \ell$ with $z_j^* = \delta^k \overline{p}_j$, that is to say j is k-headed. Then

$$\begin{aligned} z_j^* - \delta^{\ell+1}\overline{p}_j = (1-\delta^{\ell+1-k})z_j^* &\leq (1-\delta^{\ell+1-k})z^*(\mathcal{A}_i(\mathcal{T}^{(\leq k)})\backslash C) \\ &\leq z^*(\mathcal{A}_i(\mathcal{T}^{(\leq k)})\backslash C) - \delta^{\ell+2}\overline{p}(\mathcal{A}_i(\mathcal{T}^{(\leq k)})\backslash C) \\ &\leq z^*(\mathcal{A}_i(\mathcal{T})\backslash C) - \delta^{\ell+2}\overline{p}(\mathcal{A}_i(\mathcal{T})\backslash C). \end{aligned}$$

In the second inequality we use that for every $j' \in \mathcal{A}_i(\mathcal{T}^{(\leq k)})$ we have $z_{j'}^* \geq \delta^{k+1} \cdot \overline{p}_{j'}$. This implies that

$$\begin{aligned} z^*(C) &= z_j^* + z^*(C\backslash\{j\}) \\ &= z_j^* + z^*(\mathcal{A}_i(\mathcal{T}) \cap C) + z^*(C\backslash(\{j\} \cup \mathcal{A}_i(\mathcal{T}))) \\ &\leq z_j^* + z^*(\mathcal{A}_i(\mathcal{T}) \cap C) + \delta^{\ell+1}(\overline{p}(C) - \overline{p}_j - \overline{p}(\mathcal{A}_i(\mathcal{T}) \cap C)) \\ &\leq z_j^* + z^*(\mathcal{A}_i(\mathcal{T}) \cap C) + \delta^{\ell+1}(1 - \overline{p}_j - \delta\overline{p}(\mathcal{A}_i(\mathcal{T}) \cap C)) \\ &\leq z^*(\mathcal{A}_i(\mathcal{T})) + \delta^{\ell+1}(1 - \delta\overline{p}(\mathcal{A}_i(\mathcal{T}))) \leq w_i. \end{aligned}$$

□

We can now complete the proof of Theorem 3.1.

Proof of Theorem 3.1. Suppose toward contradiction there is no potential move of a k-headed job, where $k \leq K$, and no move in the blocker tree is valid. It is obvious that since Lemma 3.2-3.3 hold for (y^*, z^*), they also hold for a scaled solution $(\alpha \cdot y^*, \alpha \cdot z^*)$ with $\alpha > 0$. We can use this to obtain a solution with an arbitrarily low objective value; thereby proving that the dual is unbounded regarding makespan 1 and therefore $\mathrm{OPT}^* > 1$. □

Theorem 3.4. *The algorithm terminates in time* $\exp(\mathcal{O}(1/\epsilon \cdot \log^2(n)))$.

Proof. Let $\ell \leq K$ be the index of the last non-empty layer in $\mathcal{T}$. We will define the so-called signature vector as $s(\mathcal{T}, \sigma) = (s_1, s_2, \ldots, s_\ell)$, where s_k is given by

$$s_k = \Bigg(\sum_{(j,i,\Theta) \in \mathcal{T}_{BH}^{(k)}} [|\mathcal{J}| - |H_i|], \sum_{(j,i,\Theta) \in \mathcal{T}_{MA}^{(k)} \cup \mathcal{T}_{HA}^{(k)}} [|\mathcal{J}| - |\sigma^{-1}(i)|], \sum_{(j,i,\Theta) \in \mathcal{T}_{SA}^{(k)}} [|\mathcal{J}| - |\sigma^{-1}(i)|], \sum_{(j,i,\Theta) \in \mathcal{T}_{HL}^{(k)}} [\min M_i], \sum_{(j,i,\Theta) \in \mathcal{T}_{HM}^{(k)}} [|\mathcal{J}| - |M_i|] \Bigg).$$

Each component in s_k represents a sublayer within layer k and it is the sum over certain values associated with its blockers. Note that these values are all strictly positive, since j_{new} is not assigned and therefore $|\sigma^{-1}(i)| < |\mathcal{J}|$.

Claim 4. *(see appendix D). The signature vector increases lexicographically after polynomially many iterations of the loop.*

This means that the number of possible vectors is an upper bound on the running time (except for a polynomial factor). Each sublayer has at most $|\mathcal{J}| \cdot |\mathcal{M}|$ many blockers (since there are at most this many moves) and the value for every blocker in each of the five cases is easily bounded by $\mathcal{O}(|\mathcal{J}|)$. This implies there are at most $(\mathcal{O}(n^3))^5 = n^{\mathcal{O}(1)}$ values for each s_k. Using $K = \mathcal{O}(1/\epsilon \log(n))$ we bound the number of different signature vectors by $n^{\mathcal{O}(K)} = \exp(\mathcal{O}(1/\epsilon \log^2(n)))$. □

4 Conclusion

We have greatly improved the running time of the local search algorithm for the Restricted Assignment problem. At the same time we were able to maintain almost the same approximation ratio. We think there are two important directions for future research. The first is to improve the approximation ratio further. For this purpose, it makes sense to first find improvements for the much simpler variant of the algorithm given in [7].

The perhaps most important open question, however, is whether the running time can be brought down to a polynomial one. Recent developments in the Restricted Max-Min Fair Allocation problem indicate that a layer structure similar to the one in this paper may also help in that regard [2]. In the mentioned paper moves are only performed in large groups. This concept is referred to as laziness. The asymptotic behavior of the partition function (the number of integer partitions of a natural number) is then used in the analysis for a better bound on the number of possible signature vectors. This approach appears to have a great potential for the Restricted Assignment problem as well. In [1] it was already adapted to the special case of two processing times.

References

[1] Chidambaram Annamalai. Lazy local search meets machine scheduling. *CoRR*, abs/1611.07371, 2016.

[2] Chidambaram Annamalai, Christos Kalaitzis, and Ola Svensson. Combinatorial algorithm for restricted max-min fair allocation. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 1357–1372, 2015.

[3] Nikhil Bansal and Maxim Sviridenko. The santa claus problem. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*, pages 31–40, 2006.

[4] Deeparnab Chakrabarty, Sanjeev Khanna, and Shi Li. On $(1, \epsilon)$-restricted assignment makespan minimization. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 1087–1101, 2015.

[5] Klaus Jansen, Felix Land, and Kati Land. Bounding the running time of algorithms for scheduling and packing problems. *SIAM J. Discrete Math.*, 30(1):343–366, 2016.

[6] Klaus Jansen, Kati Land, and Marten Maack. Estimating the makespan of the two-valued restricted assignment problem. In *15th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2016, June 22-24, 2016, Reykjavik, Iceland*, pages 24:1–24:13, 2016.

[7] Klaus Jansen and Lars Rohwedder. On the configuration-lp of the restricted assignment problem. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2670–2678. SIAM, 2017.

[8] Jan Karel Lenstra, David B. Shmoys, and Éva Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Math. Program.*, 46:259–271, 1990.

[9] Lukás Polácek and Ola Svensson. Quasi-polynomial local search for restricted max-min fair allocation. *ACM Transactions on Algorithms*, 12(2):13, 2016.

[10] Ola Svensson. Santa claus schedules jobs on unrelated machines. *SIAM Journal on Computing*, 41(5):1318–1341, 2012.

A Invariants

Recall that blockers have to meet certain conditions when they are added. Variants of these conditions hold throughout the lifespan of the blocker. We will use them throughout the appendix.

Lemma A.1 (Invariants)**.** *At the start of each iteration of the loop,*

1. *for all* $(j, i, BH) = \mathcal{B} \in \mathcal{T}_{BH}^{(k)}$*:* $p(\sigma^{-1}(i) \backslash H_i) + p_j \leq 1 + R$*;*
2. *for all* $(j, i, \Theta) = \mathcal{B} \in \mathcal{T}_{HA}^{(k)} \cup \mathcal{T}_{MA}^{(k)}$*:* $p(\sigma^{-1}(i) \backslash H_i) + p_j > 1 + R$*;*
3. *for all* $(j, i, HL) = \mathcal{B} \in \mathcal{T}_{HL}^{(k)}$*:* $p(S_i(\mathcal{T}^{(\leq k)}) \cup M_i^{\min}) + p_j > 1 + R$*;*
4. *for all* $i \in \mathcal{M}(\mathcal{T}_{HL})$*:* $|M_i| \geq 1$*;*
5. *for all* $i \in \mathcal{M}(\mathcal{T}_{HM})$*:* $|M_i| \geq 2$*.*

Proof of (1). Recall that only huge jobs are undesirable on i and small or medium jobs may be moved towards i. We will proceed with the study of a move (j', i) that is to be performed and verify that the statement is true afterwards if it was before. If there is a huge job assigned to i, then $\overline{P}(\sigma^{-1}(i)) = p(\sigma^{-1}(i) \backslash H_i) + 1$. By definition of a valid move, $\overline{P}(\sigma^{-1}(i)) + p_{j'} \leq 1 + R$; hence $p(\sigma^{-1}(i) \backslash H_i) + p_{j'} \leq R$. We conclude that adding j' to $\sigma^{-1}(i)$ cannot compromise the invariant.

Conversely, suppose that no huge job is assigned to i. Then $\overline{P}(\sigma^{-1}(i)) + p_j = p(\sigma^{-1}(i) \backslash H_i) + p_j \leq 1 + R$. This means the huge-/medium-to-huge move (j, i) is valid as well. Since (j', i) is performed instead, its blocker must be in a lower sublayer. This means that $\mathcal{B}$ is deleted when (j', i) is performed and the invariant no longer needs to be shown. □

Proof of (2). Every time a job is removed from $\sigma^{-1}(i) \backslash H_i$, we need to verify that the invariant still holds. If the job was activated by $\mathcal{B}$, then the algorithm ensures this by checking the condition and deleting the blocker if it does not. Likewise, if it was activated in an earlier sublayer, $\mathcal{B}$ is deleted in any case. Finally we note that the job cannot be activated in the same sublayer by a different blocker, since no two huge-/medium-to-any blockers can coexist for the same machine. □

Proof of (3) and (4). Let $(j, i, HL) = \mathcal{B} \in \mathcal{T}_{HL}^{(k)}$ and let it be the first blocker for i added to this sublayer. First, note that $S_i(\mathcal{T}^{(\leq k)})$ has not changed since $\mathcal{B}$ was added. If a blocker was added or removed from $\mathcal{T}_{SA}^{(\leq k)} \cup \mathcal{T}_{MA}^{(\leq k)} \cup \mathcal{T}_{HA}^{(\leq k)}$, then $\mathcal{B}$ would have been deleted; hence $\mathcal{T}_{SA}^{(\leq k)} \cup \mathcal{T}_{MA}^{(\leq k)} \cup \mathcal{T}_{HA}^{(\leq k)}$ did not change. By definition, none of the jobs in $S(\mathcal{T}^{(\leq k)})$ could have been moved.

Next, we argue that at any point $p(S_i(\mathcal{T}^{(\leq k)})) + p_j \leq 1 + R < p(S_i(\mathcal{T}^{(\leq k)}) \cup M_i^{\min}) + p_j$, which implies $M_i \neq \emptyset$. This is certainly true when $\mathcal{B}$ is added to the blocker tree. The algorithm also does not move any medium jobs to i that would replace $\min M_i$ as the smallest medium job, unless it is a move of a blocker in an earlier sublayer and $\mathcal{B}$ is deleted. We therefore only need to check that the condition above still holds, when $\min M_i$ is removed from i. If $\min M_i$ was activated by $\mathcal{B}$, the algorithm ensures that the statement above holds. Otherwise, the job must have been activated by a blocker in an earlier sublayer and $\mathcal{B}$ is deleted when $\min M_i$ is removed.

Although we have only shown that (3) holds for the first blocker in the sublayer, with $M_i \neq \emptyset$ we now know that it holds for all other blockers as well ($p(S_i(\mathcal{T}^{(\leq k)})$ has not changed and $p(M_i^{\min})$ has not decreased). □

Proof of (5). Let $(j, i, HM) = \mathcal{B} \in \mathcal{T}_{HM}^{(k)}$ be the first blocker added in this sublayer for machine i. As in the previous case, $p(S_i(\mathcal{T}^{(\leq k)}))$ does not change during the lifespan of $\mathcal{B}$. Here we use $p(S_i(\mathcal{T}^{(\leq k)}) \cup M_i^{\min}) + p_j \leq 1 + R < p(S_i(\mathcal{T}^{(\leq k)}) \cup M_i) + p_j$ to show that $M_i^{\min} \neq M_i$; thus $|M_i| \geq 2$. Whenever a job is removed from M_i, it was either activated in an earlier sublayer than $\mathcal{B}$ and $\mathcal{B}$ is deleted or the algorithm verifies that the conditions still hold. Like in the previous case, the algorithm would only add jobs to M_i (and thereby possibly interfere with the statement), if it also deletes $\mathcal{B}$. □

B Unboundedness

We omitted the proof that $|\mathcal{M}(\mathcal{T}_{SA}^{(k)})| \geq |\mathcal{M}(\mathcal{T}_{HA}^{(k)})|$ for all layers k.

Proof of Claim 1. First observe the behavior of the sublayer corresponding to huge-to-any blockers in some layer k. It starts empty, then subsequently blockers are added to $\mathcal{T}_{HA}^{(k)}$. In between $\mathcal{T}_{SA}^{(k)}$ may be erased. Eventually, the huge-to-any blockers of this layer are removed all at once and the sublayer is empty again. We call this process a circle.

Let $p = |\mathcal{T}_{HA}^{(k)}|$ and $\mathcal{B}_1, \mathcal{B}_2, \ldots, \mathcal{B}_p$ be the blockers in $\mathcal{T}_{HA}^{(k)}$ in the order they were added during the current circle. Furthermore, let $\mathcal{R}_1$ ($\mathcal{R}_2, \ldots, \mathcal{R}_p$) be the blockers in $\mathcal{T}_{SA}^{(k)}$ right before $\mathcal{B}_1$ (respectively, $\mathcal{B}_2, \ldots, \mathcal{B}_p$) was added. Finally let $\mathcal{R}_{p+1} = \mathcal{T}_{SA}^{(k)}$, i.e., the current blockers. Note that not all blockers in $\mathcal{R}_1, \ldots, \mathcal{R}_{p+1}$ are necessarily still in $\mathcal{T}_{SA}^{(k)}$. We will, however, show that

$$\mathcal{M}(\mathcal{R}_{p+1}) \supsetneq \mathcal{M}(\mathcal{R}_{p+1}) \cap \mathcal{M}(\mathcal{R}_p) \supsetneq \ldots \supsetneq \mathcal{M}(\mathcal{R}_{p+1}) \cap \ldots \cap \mathcal{M}(\mathcal{R}_1).$$

This implies

$$\begin{aligned}|\mathcal{M}(\mathcal{T}_{SA}^{(k)})| = |\mathcal{M}(\mathcal{R}_{p+1})| &\geq 1 + |\mathcal{M}(\mathcal{R}_{p+1}) \cap \mathcal{M}(\mathcal{R}_p)| \\ &\geq \ldots \geq p + |\mathcal{M}(\mathcal{R}_{p+1}) \cap \ldots \cap \mathcal{M}(\mathcal{R}_1)| \geq |\mathcal{M}(\mathcal{T}_{HA}^{(k)})|.\end{aligned}$$

Let $q \in \{1, \ldots, p\}$. Suppose toward contradiction that $\mathcal{M}(\mathcal{R}_{p+1}) \cap \ldots \cap \mathcal{M}(\mathcal{R}_{q+1}) \subseteq \mathcal{M}(\mathcal{R}_q)$. Let $(j, i, \Theta) := \mathcal{B}_q$ and $\mathcal{T}'$ be the blocker tree right before B_q is added. Then in particular $\mathcal{R}_q = \mathcal{T}_{SA}'^{(k)} \subseteq \mathcal{T}'$.

Case 1. $\boldsymbol{\sigma^{-1}(i) \cap \mathcal{J}_S \supsetneq S_i^{(\leq k)}(\mathcal{T}')}$. Let j_S be a small job with $\sigma(j_S) = i$ and $j_S \notin S^{(\leq k)}(\mathcal{T}')$. By definition of $S(\mathcal{T}')$ there exists an $i' \in \Gamma(j_S) \backslash (\mathcal{M}_{BA}(\mathcal{T}') \cup \mathcal{M}_{MA}(\mathcal{T}') \cup \mathcal{M}_{SA}(\mathcal{T}') \cup \{i\})$. This implies $i' \notin \mathcal{R}_q$. However, i' must be in $\mathcal{M}(\mathcal{R}_{q+1})$, because otherwise (j_S, i') would be a potential move with a higher priority than $\mathcal{B}_{q+1}$ at the time it was added. For the same reason i' must be in $\mathcal{M}(\mathcal{R}_{q+2}), \ldots, \mathcal{M}(\mathcal{R}_{p+1})$. This is a contradiction to $\mathcal{M}(\mathcal{R}_{p+1}) \cap \ldots \cap \mathcal{M}(\mathcal{R}_{q+1}) \subseteq \mathcal{M}(\mathcal{R}_q)$.

Case 2. $\boldsymbol{\sigma^{-1}(i) \cap \mathcal{J}_S = S_i^{(\leq k)}(\mathcal{T}')}$. Since no jobs were moved to i, neither $p(M_i)$, nor $p(S_i^{(\leq k)}(\mathcal{T}'))$ have increased since $\mathcal{B}_q$ was added. Therefore we still have $p(S_i^{(\leq k)}(\mathcal{T}') \cup M_i) + p_j \leq 1 + R$. This implies

$$p(\sigma^{-1}(i) \backslash H_i) + p_j = p(S_i^{(\leq k)}(\mathcal{T}') \cup M_i) + p_j \leq 1 + R\,,$$

a contradiction to Invariant 2. □

C Feasibility

Proof of Claim 2. Let $k \leq K$, $i \notin \mathcal{M}(\mathcal{T}_{SA}^{(\leq k)} \cup \mathcal{T}_{MA}^{(\leq k)} \cup \mathcal{T}_{HA}^{(\leq k)})$, $C \in \mathcal{C}_i$, $j \in C$ k-headed and big with $\sigma(j) \neq i$. Recall we have to show that $z_j^* \leq z^*(\mathcal{A}_i(\mathcal{T}^{(\leq k)}) \backslash C)$. Consider the types of blockers regarding i.

Case 1. $\boldsymbol{i \notin \mathcal{M}(\mathcal{T}^{(\leq k)})}$. Then since (j, i) is not a potential move, we find that j must be huge and $p(S_i(\mathcal{T}^{(\leq k)})) + p_j > 1 + R$. This implies

$$\begin{aligned}z^*(\mathcal{A}_i(\mathcal{T}^{(\leq k)}) \backslash C) &\geq z^*(S_i(\mathcal{T}^{(\leq k)}) \backslash C) \\ &\geq \delta^k(p(S_i(\mathcal{T}^{(\leq k)})) - p(S_i(\mathcal{T}^{(\leq k)}) \cap C)) \\ &\geq \delta^k(p(S_i(\mathcal{T}^{(\leq k)})) - (1 - p_j)) > \delta^k \cdot R > z_j^*.\end{aligned}$$

Case 2. $\boldsymbol{i \in \mathcal{M}(\mathcal{T}_{BH}^{(\leq k)})}$. We argue that there has to be a huge job on i. Let j_B be the big job corresponding to the huge-/medium-to-huge blocker for i. If $H_i = \emptyset$ then using Invariant 1 we have

$$\overline{P}(\sigma^{-1}(i)) + p_{j_B} = p(\sigma^{-1}(i) \backslash H_i) + p_{j_B} \leq 1 + R\,;$$

a contradiction, since (j_B, i) is not valid. We conclude that there has to be a huge job $j_H \in \mathcal{A}_i(\mathcal{T}^{(\leq k)})$ that is still assigned to i. This implies that

$$z^*_j \leq \delta^k \frac{5}{6} \leq z^*_{j_H} \leq z^*(\mathcal{A}_i(\mathcal{T}^{(\leq k)})\backslash C).$$

Case 3. $\boldsymbol{i \in \mathcal{M}(\mathcal{T}^{(\leq k)}_{HM})}$**.** There are at least two jobs j_M and j'_M in M_i (Invariant 5). Using $\epsilon < 1/12$ we get

$$z^*_j \leq \delta^k \frac{5}{6} \leq \delta^{k+1}(\overline{p}_{j_M} + \overline{p}_{j'_M}) \leq z^*_{j_M} + z^*_{j'_M} \leq z^*(\mathcal{A}_i(\mathcal{T}^{(\leq k)})\backslash C).$$

Case 4. $\boldsymbol{i \in \mathcal{M}(\mathcal{T}^{(\leq k)}_{HL})\backslash(\mathcal{M}(\mathcal{T}^{(\leq k)}_{BH}) \cup \mathcal{M}(\mathcal{T}^{(\leq k)}_{HM}))}$**.** By Invariant 4 we know that M_i is not empty. Set $j_M = \min M_i$.

Case 4.1. $j > j_M$ or (j, i) is in $\mathcal{T}^{(\leq k)}$. If (j, i) is not in $\mathcal{T}^{(\leq k)}$, then because it is not a potential move and j is not undesirable on i, we have $1 + R < p(S_i(\mathcal{T}^{(\leq k)})) + p_j \leq p(S_i(\mathcal{T}^{(\leq k)}) \cup \{j_M\}) + p_j$. If it is a blocker, then in particular $p(S_i(\mathcal{T}^{(\leq k)}) \cup \{j_M\}) + p_j > 1 + R$, because it must be a huge-to-least blocker (Invariant 3). Thus

$$\begin{aligned} z^*(\mathcal{A}_i(\mathcal{T}^{(\leq k)})\backslash C) &\geq \delta^k p(S_i(\mathcal{T}^{(\leq k)})\backslash C) + \delta^{k+1} \cdot p_{j_M} \\ &\geq \delta^k p(S_i(\mathcal{T}^{(\leq k)})) - \delta^k(1 - p_j) + \delta^k p_{j_M} - \epsilon\delta^k p_{j_M} \\ &> \delta^k \left(\frac{11}{6} + \epsilon - p_j\right) - \delta^k(1 - p_j + \epsilon) \geq \delta^k \frac{5}{6} \geq z^*_j. \end{aligned}$$

Case 4.2. $j < j_M$ and (j, i) is not a blocker in $\mathcal{T}^{(\leq k)}$. Then in particular j is medium. We now want to show that j_M is not $(k+1)$-headed and thus $z^*_{j_M} \geq \delta^k \overline{p}_{j_M} \geq \delta^k \overline{p}_j \geq z^*_j$. Since j is k-headed and medium, it must be activated by a blocker in layer $k-1$. Suppose toward contradiction that $i \notin \mathcal{M}(\mathcal{T}^{(\leq k-1)}_{HL})$. At the time the huge-to-least blocker was added, j was already activated. Since no jobs were undesirable on i, the move (j, i) must have been a potential move as well. Moves of medium jobs have a higher priority than huge-to-least moves and (j, i) would have been chosen instead.

We conclude that $i \in \mathcal{M}(\mathcal{T}^{(k-1)}_{HL})$, j_M is k-headed, and

$$z^*_j \leq z^*_{j_M} \leq z^*(\mathcal{A}_i(\mathcal{T}^{(\leq k)})\backslash C). \qquad \square$$

Proof of Claim 3. Let $k \leq K$ and $i \in \mathcal{M}(\mathcal{T}^{(k)}_{SA} \cup \mathcal{T}^{(k)}_{MA} \cup \mathcal{T}^{(k)}_{HA})$. We claim that $w_i \geq z^*(\mathcal{A}_i(\mathcal{T})) + \delta^k(1 - \delta\overline{p}(\mathcal{A}_i(\mathcal{T})))$. Recall that all jobs on i are activated.

Case 1. $\boldsymbol{i \in \mathcal{M}(\mathcal{T}^{(k)}_{HA})}$**.** Then there must be a huge-to-any move (j_B, i) in $\mathcal{T}$. This implies $1 + 2\epsilon - 1/6 \leq 1 + R - p_{j_B} < p(\sigma^{-1}(i)\backslash H_i) \leq \overline{p}(\mathcal{A}_i(\mathcal{T}))$ (Invariant 2). Since $\overline{p}(\mathcal{A}_i(\mathcal{T})) \leq 1 + R < 2$, we find

$$\begin{aligned} w_i = z^*(\mathcal{A}_i(\mathcal{T})) + \delta^k \frac{1}{6} &> z^*(\mathcal{A}_i(\mathcal{T})) + \delta^k(1 + 2\epsilon - \overline{p}(\mathcal{A}_i(\mathcal{T}))) \\ &\geq z^*(\mathcal{A}_i(\mathcal{T})) + \delta^k(1 - \delta\overline{p}(\mathcal{A}_i(\mathcal{T}))). \end{aligned}$$

Case 2. $\boldsymbol{i \in \mathcal{M}(\mathcal{T}^{(k)}_{MA})}$**.** Let (j_M, i) be the move corresponding to the medium-to-any blocker. Then $1 + 2\epsilon \leq 1 + R - p_{j_M} < p(\sigma^{-1}(i)\backslash H_i) \leq \overline{p}(\mathcal{A}_i(\mathcal{T}))$ (Invariant 2) and

$$\begin{aligned} w_i = z^*(\mathcal{A}_i(\mathcal{T})) &> z^*(\mathcal{A}_i(\mathcal{T})) + \delta^k(1 + 2\epsilon - \overline{p}(\mathcal{A}_i(\mathcal{T}))) \\ &\geq z^*(\mathcal{A}_i(\mathcal{T})) + \delta^k(1 - \delta\overline{p}(\mathcal{A}_i(\mathcal{T}))). \end{aligned}$$

Case 3. $\boldsymbol{i \in \mathcal{M}(\mathcal{T}^{(k)}_{SA})}$**.** Then there is a small move (j_S, i) in $\mathcal{T}$, which is not valid. Since i is assigned at most one huge job, we have

$$\overline{p}(\sigma^{-1}(i)) \geq \overline{P}(\sigma^{-1}(i)) - \frac{1}{6} > 1 + R - p_{j_S} - \frac{1}{6} \geq 1 + \frac{1}{6} + 2\epsilon.$$

We conclude

$$w_i = z^*(\mathcal{A}_i(\mathcal{T})) - \delta^k \frac{1}{6} > z^*(\mathcal{A}_i(\mathcal{T})) + \delta^k(1 + 2\epsilon - \overline{p}(\mathcal{A}_i(\mathcal{T})))$$
$$\geq z^*(\mathcal{A}_i(\mathcal{T})) + \delta^k(1 - \delta\overline{p}(\mathcal{A}_i(\mathcal{T}))).$$

□

D Termination

Proof of Claim 4. When the algorithm adds a blocker, all sublayers after this blocker will be deleted. However, the contribution of the new blocker to its sublayer will overall increase the lexicographic value of the signature vector.

Now consider a valid move that is performed. We will observe the last move (j_0, i_0) of potentially many consecutive moves and show that afterwards the lexicographic value of the signature vector has increased. Let $\mathcal{B} = (j, i, \Theta)$ be the blocker that activated j_0. In particular, j_0 was assigned to i. There are two cases to consider. In the first one, $\mathcal{B}$ and its sublayer remain unchanged and all sublayers after $\mathcal{B}$ are removed. Then the component corresponding to the sublayer of $\mathcal{B}$ increases, because j_0 is removed from i: In all except the huge-to-least blockers, the value of the blocker is essentially the negation of the number of undesirable jobs on the machine. During the lifetime of a blocker this number does not decrease. When we remove j_0, it even strictly increases. It is easy to see that for huge-to-least blockers the value, that is $\min M_i$, increases as well, since the smallest medium job is removed.

The critical case is, when $\mathcal{B}$ and its sublayer are removed as well (see algorithm, lines 13-15). The algorithm will add a new blocker in the next iteration. We will show that this blocker will be added to a smaller sublayer than $\mathcal{B}$ was, thereby ultimately increasing the lexicographic order. We will show that (i) (j, i) is a potential move to a smaller sublayer and (ii) every move with a higher or the same priority as (j, i) is in a smaller sublayer as well.

For (i) we figure that compared to the last time (j, i) was added as a blocker, no new blockers are in $\mathcal{T}$. In particular j is not undesirable on i. If a blocker was added to a lower sublayer than $\mathcal{B}$ was, then $\mathcal{B}$ would already have been removed. On the other hand, the sublayer of $\mathcal{B}$ and all later sublayers are deleted, so they do not contain new blockers either. We understand that j cannot be undesirable on i, since it was not undesirable when blocker $\mathcal{B}$ was added.

Case 1. **$\mathcal{B}$ was a medium-to-any or huge-to-any blocker.** Then we now have $p(\sigma^{-1}(i) \setminus H_i) + p_j \leq 1 + R$ and (j, i) is a potential huge-/medium-to-huge blocker. Any move with an higher or equal priority to (j, i) must be either a huge-/medium-to-huge blocker in the same layer or a blocker in an earlier layer. In particular it must be in an earlier sublayer than $\mathcal{B}$ was.

Case 2. **$\mathcal{B}$ was a huge-to-least blocker.** We argue that M_i must be empty: Suppose toward contradiction that $M_i \neq \emptyset$. Invariant 3 tells us, that before j_0 was moved, we had $p(S_i(\mathcal{T}^{(\leq k)}) \cup M_i^{\min}) + p_j > 1 + R$. The move has no effect on $S_i(\mathcal{T}^{(\leq k)})$ and since M_i is not empty, $p(M_i^{\min})$ can only increase; a contradiction.

$M_i = \emptyset$ implies that

$$p(S_i(\mathcal{T}^{(\leq k)}) \cup M_i) + p_j = p(S_i(\mathcal{T}^{(\leq k)}) \cup M_i^{\min}) + p_j \leq 1 + R;$$

hence (j, i) is a huge-to-any or huge-to-huge move. A move with a higher or equal priority in the same layer must be a huge-/medium-to-huge, huge-to-any, or small-to-any move, all of which are in earlier sublayers.

Case 3. **$\mathcal{B}$ was a huge-to-medium blocker.** Then one of the two conditions for huge-to-medium blockers no longer holds. $p(S_i(\mathcal{T}^{(\leq k)}))$ has not increased; hence it must still be some potential move. Any blocker in the same layer with higher or equal priority than (j, i) cannot be a huge-to-medium blocker either and is therefore in a lower sublayer.

□